

Approximately Truthful Multi-Agent Optimization Using Cloud-Enforced Joint Differential Privacy

M.T. Hale and M. Egerstedt*

Abstract

Multi-agent coordination problems often require agents to exchange state information in order to reach some collective goal, such as agreement on a final state value. In some cases, it is feasible that opportunistic agents may deceptively report false state values for their own benefit, e.g., to claim a larger portion of shared resources. Motivated by such cases, this paper presents a multi-agent coordination framework which disincentivizes opportunistic misreporting of state information. This paper focuses on multi-agent coordination problems that can be stated as nonlinear programs, with non-separable constraints coupling the agents. In this setting, an opportunistic agent may be tempted to skew the problem's constraints in its favor to reduce its local cost, and this is exactly the behavior we seek to disincentivize. The framework presented uses a primal-dual approach wherein the agents compute primal updates and a centralized cloud computer computes dual updates. All computations performed by the cloud are carried out in a way that enforces *joint differential privacy*, which adds noise in order to dilute any agent's influence upon the value of its cost function in the problem. We show that this dilution deters agents from intentionally misreporting their states to the cloud, and present bounds on the possible cost reduction an agent can attain through misreporting its state. This work extends our earlier work on incorporating ordinary differential privacy into multi-agent optimization, and we show that this work can be modified to provide a disincentive for misreporting states to the cloud. Numerical results are presented to demonstrate convergence of the optimization algorithm under joint differential privacy.

I. INTRODUCTION

Multi-agent optimization problems have found applications in a range of research areas, including power systems in [17], machine learning in [3], sensor networks in [5], and robotics in [19]. Solutions to some problems in these applications rely on the implicit assumption that all agents share correct, truthful information with others in a network. However, one can envision a scenario in which deceitful agents do not honestly share their data, instead sharing false information that will skew the behavior of the system in their favor. For example, a homeowner connected to a smart power grid may report a false value for his or her power usage in order to save money. This paper considers optimization problems with agents that may be intentionally deceitful for their own benefit, and it provides a method for disincentivizing untruthful behavior when teams of agents are collectively optimizing.

We reduce the incentive to share false information by using *joint differential privacy*, defined in [13], to limit the possible decrease in cost an agent can achieve through intentionally misreporting its state. Joint differential privacy

*The authors are with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA. Email: {matthale, magnus}@gatech.edu. Research supported in part by the NSF under Grant CNS-1239225.

adds noise to reduce the ability of an agent to benefit from reporting false information, and was first introduced for this purpose in [13] to promote truthful sharing of information by the players in a class of games. This work was applied specifically to optimization problems arising from distributed electric vehicle charging in [11], as well as to linearly separable optimization problems in [12].

The developments in this paper differ from those in [11] and [12], as well as some other work on private optimization, e.g., [10], in two key ways. First, the work in [11], [12], and [10] uses differential privacy as it was originally defined in [7] to keep some static object, such as a constraint function, private. Our work focuses on trajectory-level privacy for state trajectories that are dynamically generated as an optimization algorithm is executed. In applications such as phase synchronization in smart power grids, each iterate of an optimization algorithm corresponds to a physical state at some point in time, and an agent’s contribution to the optimization process is its whole state trajectory. Applying joint differential privacy in such applications should therefore be done at the trajectory level, and implementing joint differential privacy for trajectory-level data is most naturally done using the dynamical systems formulation of differential privacy. Therefore, this paper uses trajectory-level privacy as defined in [15], rather than privacy for databases as in [7].

This form of privacy is not only better suited to the applications of interest, but also allows privacy guarantees to hold across infinite time horizons which, in many cases, cannot be attained using privacy for databases. To our knowledge, this paper is the first use of joint differential privacy at the trajectory level. The second key difference between the existing literature and this paper is the class of optimization problems solved. We apply joint differential privacy in a general multi-agent nonlinear programming setting, incorporating non-separable functional constraints that are also possibly nonlinear, which differs from work done in [11], [12], and [10] where either linear or affine constraints are considered.

To solve such problems, a cloud computer is added to the team of agents to serve as a trusted central aggregator, also sometimes called a “curator” in the privacy literature. This architecture was previously used for privately solving multi-agent nonlinear programs in [9], in which honest agents seek to protect sensitive data from eavesdroppers via ordinary differential privacy. The work in [16] suggests that any differential privacy implementation provides some disincentive against untruthful information sharing, and the current paper uses the notion of joint differential privacy to formally provide this guarantee in the framework developed in [9]. Accordingly, the technical novelty of this paper is not in the algorithm used to solve problems, but in the theoretical performance guarantees that are provided using this framework; the work in [9] focuses exclusively on protecting sensitive data, but the current paper focuses on the problem of preventing untruthful behavior by the agents. The technical contribution of this paper thus consists of adapting our existing privacy framework to the problem of incentivizing truthful behavior by the agents, and quantifying the extent to which any agent can benefit from untruthful behavior.

Several existing approaches use behavioral analysis to identify untruthful agents, including those in [4] and [8]. In this paper, an agent’s local state updates rely on a local objective function and local constraint set that are considered sensitive, and therefore these local data are not shared with any other agent. As a result, the correct next value of an agent’s state is only known to that agent, and a behavioral analysis approach cannot be used here because no

outside observer can determine what any agent's future states should be. Rather than detecting untruthful behavior, this paper seeks to prevent it outright by reducing the incentive to share untruthful information via joint differential privacy. The principle underlying this application of joint differential privacy is that adding noise to the system can make each agent's cost insensitive to changes in the agent's state trajectory. The amount of noise added must be calibrated to the system in order to dilute the effect of an agent reporting an untruthful state to the cloud, and this paper presents this calibration for a general multi-agent nonlinear program in terms of constants pertaining to the problem.

The problems considered consist of a collection of agents, each with a local objective function and local set constraint, and ensemble state constraints that jointly constrain the agents. A primal-dual approach is used, in which the agents update their own states, which are the problem's primal variables, and the cloud updates the problem's dual variables. Naturally, the constraints in this problem will usually lead to higher costs for each agent than a comparable unconstrained problem. As a result, some of the agents may wish to skew the constraints in their favor. One way they may do so is by reporting false state information to the cloud in order to loosen the constraints' effects on their own states, thereby giving the untruthful agents a lower cost. This manipulation of the constraints will result in an unequal distribution of the burden of these constraints by tightening them on honest agents. Therefore, this form of untruthful behavior is disincentivized using joint differential privacy.

The rest of the paper is organized as follows. Section II provides the necessary optimization background, the structure of communications in the system, and formally states the joint differentially private optimization problem that is the focus of the paper. Then, Section III reviews joint differential privacy, and Section IV presents the proposed joint differentially private algorithm. Next, Section V proves the main result of the paper on limiting an agent's incentive to misreport its states to the cloud. Section VI then presents simulation results and Section VII concludes the paper.

II. BACKGROUND ON OPTIMIZATION AND PROBLEM STATEMENT

This section presents the multi-agent optimization problem of interest and an optimization algorithm that will later be used to solve a joint differentially private version of this problem. This section also describes the cloud-based architecture used to solve this problem. Throughout the paper, ∇_{x_i} denotes the partial derivative with respect to x_i , and h_{x_i} denotes the partial derivative of the function h with respect to x_i .

A. Optimization Problem Formulation

Consider a problem consisting of N agents indexed over $i \in [N] := \{1, \dots, N\}$. Agent i has state $x_i \in \mathbb{R}^{n_i}$ with $n_i \in \mathbb{N}$ and a local set constraint of the form $x_i \in X_i \subset \mathbb{R}^{n_i}$. The diameter of X_i is denoted by $D_i := \max_{x_1, x_2 \in X_i} \|x_1 - x_2\|_1$. Agent i also has a local objective function $f_i : \mathbb{R}^{n_i} \rightarrow \mathbb{R}$ depending only upon its own state. Agent i 's local data are subject to the following assumption.

Assumption 1: For all $i \in [N]$, f_i is C^2 and convex in x_i , and X_i is non-empty, compact, and convex. \triangle

Assumption 1 implies that f_i is Lipschitz and its Lipschitz¹ constant is denoted by K_i . In particular, Assumption 1 allows for convex polynomial objectives and box constraints, which are common in multi-agent optimization problems. Both f_i and X_i are considered sensitive information and are therefore not shared with the other agents or with the cloud. For simplicity of notation, define the set $X := X_1 \times \cdots \times X_N \subset \mathbb{R}^n$, where $n = \sum_{i \in [N]} n_i$. The agents' individual set constraints require $x \in X$, where x is the ensemble state vector of the network, defined as $x = (x_1^T, \dots, x_N^T)^T \in X$, and where Assumption 1 provides that X is non-empty, compact, and convex.

The agents are together subject to global inequality constraints $g(x) \leq 0$, where $g : \mathbb{R}^n \rightarrow \mathbb{R}^m$. The functional constraints in g are subject to the following assumption.

Assumption 2:

- i. For all $j \in [m]$, the constraint g_j is C^2 and convex in x .
- ii. There exists a point $\bar{x} \in X$ such that $g(\bar{x}) < 0$. \triangle

Assumption 2.i. admits a wide variety of constraint functions, e.g., any convex polynomials. Assumption 2.ii. is known as Slater's condition, e.g., Assumption 6.4.2 in [2], and, in conjunction with Assumption 2.i., guarantees that strong duality holds. Assumption 2 implies that g and each g_{x_i} are Lipschitz. Their Lipschitz constants are denoted by K_g and $L_{g,i}$, respectively.

Summing the per-agent objective functions gives the ensemble objective $f(x) = \sum_{i \in [N]} f_i(x_i)$, which is C^2 and convex in x because of Assumption 1. Together f , g , and X comprise the following ensemble-level optimization problem.

Problem 0: (Preliminary; no joint differential privacy yet)

$$\begin{aligned} & \text{minimize } f(x) \\ & \text{subject to } g(x) \leq 0 \\ & x \in X. \end{aligned} \quad \diamond$$

We now detail how to solve Problem 0 in the cloud-based system, and then give a unified problem statement, including incentivizing truthful behavior, in Problem 1 below.

The Lagrangian associated with Problem 0 is

$$L(x, \mu) = f(x) + \mu^T g(x),$$

where $\mu \in \mathbb{R}_+^m$ is the dual vector associated with Problem 0 and \mathbb{R}_+^m denotes the non-negative orthant of \mathbb{R}^m . Seminal work of [14] shows that a point $\hat{x} \in X$ is a solution to Problem 0 if and only if there exists a point $\hat{\mu} \in \mathbb{R}_+^m$ such that $(\hat{x}, \hat{\mu})$ is a saddle point of L . This saddle point condition can be compactly expressed by requiring

$$L(\hat{x}, \mu) \leq L(\hat{x}, \hat{\mu}) \leq L(x, \hat{\mu}) \text{ for all } (x, \mu) \in X \times \mathbb{R}_+^m, \quad (1)$$

and an optimal primal-dual pair $(\hat{x}, \hat{\mu})$ exists under Assumptions 1 and 2.

¹All Lipschitz constants in this paper are with respect to the metric induced by the 1-norm.

For the forthcoming optimization algorithm, Equation (II-A) is used to define an upper bound on the norm of $\hat{\mu}$, stated in the following lemma. It uses the Slater point \bar{x} from Assumption 2.ii. and any lower bound on f over X , denoted f_{lower} , which exists under Assumption 1.

Lemma 1: For $(\hat{x}, \hat{\mu})$ a saddle point of L ,

$$\hat{\mu} \in M := \left\{ \mu \in \mathbb{R}_+^m : \|\mu\|_1 \leq \frac{f(\bar{x}) - f_{lower}}{\min_{j \in [m]} \{-g_j(\bar{x})\}} \right\}.$$

Proof: See Section II-A in [9]. ■

Using the fact that a saddle point of L provides a solution to Problem 0, the remainder of the paper focuses on finding such saddle points. The algorithm used for this purpose includes an iterative Tikhonov regularization with an asymptotically vanishing stepsize, and was stated in [1] for deterministic variational inequalities and later in [18] for stochastic problems. To help describe the communications and computations in the cloud-based system, we provide the general deterministic form of this algorithm now, though in Section IV stochasticity is introduced when the algorithm is made joint differentially private. The saddle-point finding algorithm uses the coupled update equations

$$x(k+1) = \Pi_X [x(k) - \gamma_k (L_x(k) + \alpha_k x(k))] \quad (2a)$$

$$\mu(k+1) = \Pi_M [\mu(k) + \gamma_k (L_\mu(k) - \alpha_k \mu(k))], \quad (2b)$$

where Π_X and Π_M are the Euclidean projections onto X and M , respectively. This update law will be referred to as Update (1). Here, γ_k is a stepsize and α_k is the regularization parameter, and the values of both will be provided in Theorem 1 in Section IV.

B. Optimizing via the Cloud

We now elaborate on the cloud-based architecture to be used and the means of executing Update (1) on the cloud and agents. To enforce joint differential privacy, the agents do not directly share any information with each other. Instead, the agents route messages through a trusted cloud computer which aggregates all states in the network, performs computations involving these states, and makes the results of these computations joint differentially private before sending them to the agents.

Splitting Equation (1) into each agent's states gives the update

$$x_i(k+1) = \Pi_{X_i} [x_i(k) - \gamma_k (L_{x_i}(k) + \alpha_k x_i(k))]$$

for agent i , where expanding the term containing L_{x_i} gives

$$x_i(k+1) = \Pi_{X_i} \left[x_i(k) - \gamma_k \left(\nabla_{x_i} f_i(x_i(k)) + g_{x_i}(x(k))^T \mu(k) + \alpha_k x_i(k) \right) \right]. \quad (3)$$

Importantly, the right-hand side of Equation (II-B) contains two terms which agent i cannot compute on its own: $g_{x_i}(x(k))$, because it is a function of every state in the network, and $\mu(k)$, because its update law relies on $g(x(k-1))$, which is also a function of every state in the network. For this reason, the cloud computer is used to compute $\mu(k)$ and $g_{x_i}(x(k))$ for each $i \in [N]$ at all timesteps k .

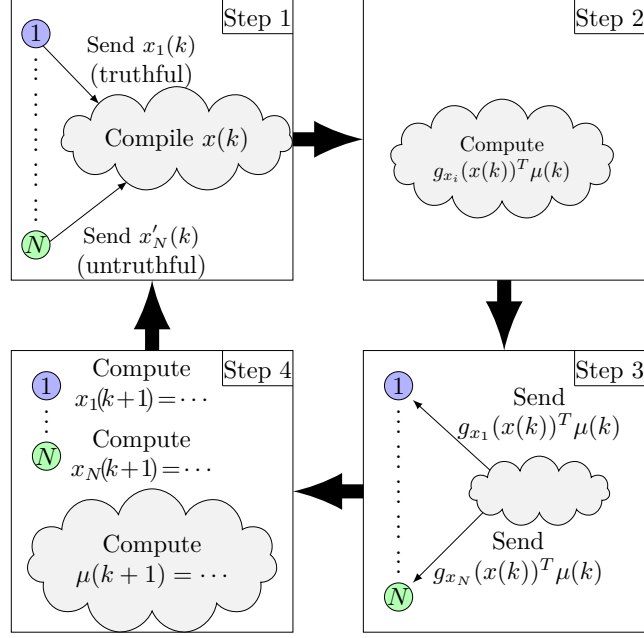


Fig. 1. The four steps of a communications cycle in the cloud-based system. First, each agent sends its state to the cloud. Second, the cloud performs centralized computations required by the agents. Third, the cloud sends the results of these computations to the agents. Fourth, agent i computes $x_i(k+1)$ and the cloud computes $\mu(k+1)$ and this process repeats. In Step 1 depicted here, agent N is misreporting its state to the cloud by sending $x'_N(k)$ instead of $x_N(k)$. As the algorithm progresses this untruthful state propagates through the system.

Four actions occur within timestep k . First, agent i sends $x_i(k)$ to the cloud and the cloud assembles the ensemble state vector $x(k)$. Second, the cloud computes $g_{x_i}(x(k))$ for every $i \in [N]$. Third, the cloud sends $g_{x_i}(x(k))^T \mu(k)$ to agent i . Fourth, the cloud computes $\mu(k+1)$ and simultaneously agent i computes $x_i(k+1)$, and then this process repeats. This exchange of information is depicted in Figure 1. In the upper-left panel of Figure 1, agents 1 through $N-1$ report their actual state to the cloud, while agent N misreports its state to the cloud by sending some $x'_N(k)$ instead of $x_N(k)$. The arrows connecting boxes indicate how misreported states propagate through the system, eventually affecting all agents' states. The cloud's computations in Steps 2 and 4 will be modified in Section III to introduce joint differential privacy, though the overall communications structure will remain the same.

To reflect that agent i receives the vector $g_{x_i}(x(k))^T \mu(k)$ from the cloud without knowing $g_{x_i}(x(k))$ or $\mu(k)$ individually, agent i 's update law is rewritten as

$$x_i(k+1) = \Pi_{X_i} \left[x_i(k) - \gamma_k (\nabla_i f_i(x_i(k)) + q_i(k) + \alpha_k x_i(k)) \right],$$

where $\mathbb{R}^{n_i} \ni q_i(k) := g_{x_i}(x(k))^T \mu(k)$. As in Equation (1), the cloud computes the dual update according to

$$\mu(k+1) = \Pi_M [\mu(k) + \gamma_k (L_\mu(k) - \alpha_k \mu(k))],$$

C. Full Problem Statement

We now state the algorithm that will be used throughout the remainder of the paper. Below that, we identify the potential for misreporting states in this algorithm and state the problem that is later solved using joint differential privacy.

Algorithm 1:

Step 0: For all $i \in [N]$, initialize agent i with $x_i(0)$, X_i , f_i , $\{\alpha_k\}_{k \in \mathbb{N}}$, and $\{\gamma_k\}_{k \in \mathbb{N}}$. Initialize the cloud with \bar{x} , g , f_{lower} , $\{\alpha_k\}_{k \in \mathbb{N}}$, and $\{\gamma_k\}_{k \in \mathbb{N}}$. Let the cloud compute M before the system begins optimizing. Set $k = 0$.

Step 1: For all $i \in [N]$, the cloud computes $q_i(k)$ and sends it to agent i .

Step 2: Agent i computes

$$x_i(k+1) = \Pi_{X_i} \left[x_i(k) - \gamma_k (\nabla_i f_i(x_i(k)) + q_i(k) + \alpha_k x_i(k)) \right],$$

and sends a state value $x_i(k+1)$ to the cloud.

Step 3: The cloud computes

$$\mu(k+1) = \Pi_M [\mu(k) + \gamma_k (L_\mu(k) - \alpha_k \mu(k))].$$

Step 4: Set $k := k+1$ and return to Step 1. ◇

In Step 2 of Algorithm 1, agent i may report a false state value to the cloud, sending some $\tilde{x}_i(k+1) \neq x_i(k+1)$. This is the misreporting behavior that this paper seeks to prevent using the cloud. The only influence the cloud has upon agent i is through $q_i(k)$, and therefore the cloud must compute $q_i(k)$ in a way that incentivizes agent i to honestly report its state. The incentivization of this behavior is stated as Problem 1 below, and the remainder of the paper focuses on solving Problem 1.

Problem 1: Execute Algorithm 1 with the cloud computing $q_i(k)$ in a way that incentivizes agent i to honestly report $x_i(k+1)$ in Step 2, while still converging to a minimum. ◇

It is assumed that each agent ultimately wants the constraints g to be satisfied, as would be the case when g corresponds to some mission-critical conditions that must be satisfied by the agents. However, an agent may wish to reduce the impact g has upon its own state in order to reduce its cost. One way of affecting g for this purpose is by reporting false states to the cloud over time; because all agents want g to be satisfied, a misreporting agent will still use $q_i(k)$ from the cloud in its state updates, but an agent can substantially influence these messages for its own benefit through misreport. In response to misreported states, other agents' messages from the cloud in Update (1) will be affected in a way that compensates for agent i 's false reported states, thereby resulting in an unfair distribution of the burden of g .

This behavior cannot be detected by the other agents or the cloud because only agent i knows f_i and X_i , making no other entity in the network capable of determining what agent i 's state should be (cf. Equation (1)). Therefore, rather than detecting manipulation of g , we seek to prevent this behavior. Joint differential privacy provides a framework for incentivizing truthful sharing of information and it is used here to prevent the agents

from manipulating g .

Because agents are opportunistic but not malicious, they may send untruthful states to the cloud, but they will not send states that harm the system or prevent convergence of Algorithm 1. As a result, a misreported state trajectory will have some relationship to an agent's true state trajectory and this relationship is used in defining adjacency of signals in our joint differential privacy implementation. The next section details the manner in which noise is added to $q_i(k)$ using joint differential privacy, and Section IV shows that this noise still allows for Algorithm 1 to reach a minimum.

III. JOINT DIFFERENTIAL PRIVACY

This section recalls necessary details from both ordinary differential privacy and joint differential privacy. Then it presents the joint differential privacy mechanism that will be implemented on the preceding cloud architecture. It is critically important to note that while this section discusses privacy, the ultimate goal is to apply joint differential privacy to induce truthful sharing of states. All of the content on privacy in this section should therefore be understood as making progress toward inducing truthful behavior.

A. Differential Privacy Background

Differential privacy as described in [7] was originally designed to keep individual database entries private whenever a database is queried, and this is done by adding noise to the responses to such queries. This idea was extended to dynamical systems in [15] in order to keep inputs to a system private from anyone observing the outputs of that system. It is the dynamical systems notion of differential privacy that is used below.

The key idea behind differential privacy is that noise is added to make “adjacent” inputs produce “similar” outputs, and these notions are made rigorous below. As above, let there be N users, with the i^{th} user contributing a signal $u_i \in \tilde{\ell}_{p_i}^{s_i}$. The space $\tilde{\ell}_{p_i}^{s_i}$ is the space of sequences of s_i -vectors in which every finite truncation of every element has finite p_i -norm. More explicitly, with $u_i(k) \in \mathbb{R}^{s_i}$ denoting the k^{th} element of the signal u_i , define the truncation operator P_t according to

$$P_t u = \begin{cases} u(k) & k \leq t \\ 0 & k > t \end{cases}.$$

Then $u_i \in \tilde{\ell}_{p_i}^{s_i}$ if and only if $P_t u_i \in \ell_{p_i}^{s_i}$ for all $t \in \mathbb{N}$. The full input space is then defined by the Cartesian product $\tilde{\ell}_p^s = \prod_{i=1}^n \tilde{\ell}_{p_i}^{s_i}$. This paper focuses on the case where $p_i = 1$ for all $i \in [N]$.

To formalize the notion of adjacency of inputs in $\tilde{\ell}_p^s$, we define a binary, symmetric adjacency relation Adj_B^i , which is parameterized by $B > 0$ and an index $i \in [N]$. Its definition uses the notation $u_{-i} = (u_1, \dots, u_{i-1}, u_{i+1}, \dots, u_N)$. The adjacency relation takes the form $\text{Adj}_B^i : \tilde{\ell}_p^s \times \tilde{\ell}_p^s \rightarrow \{0, 1\}$ and has the following definition, as stated in [15].

Definition 1: Two inputs $u, \tilde{u} \in \tilde{\ell}_p^s$ satisfy $\text{Adj}_B^i(u, \tilde{u}) = 1$ for some $i \in [N]$ if and only if $\|u_i - \tilde{u}_i\|_{p_i} \leq B$ and $u_{-i} = \tilde{u}_{-i}$. The signals u and \tilde{u} are then said to be adjacent with respect to agent i . If i is arbitrary, the relation Adj_B is used, and u and \tilde{u} are simply called adjacent. ♦

This section maintains use of the symbol u for system inputs (rather than x as will be in subsequent sections) to maintain continuity with the references cited here for private dynamical systems.

Inputs from $\tilde{\ell}_p^s$ are assumed to pass into a causal, deterministic system \mathcal{G} , which produces outputs in $\tilde{\ell}_q^r$. To define when two outputs are “similar,” the notion of a *mechanism* is used. In the context of private dynamical systems, a mechanism is a means of adding noise to an otherwise deterministic system to provide privacy to that system’s inputs. Formally, for a fixed a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, a mechanism is a map of the form

$$M : \tilde{\ell}_p^s \times \Omega \rightarrow \tilde{\ell}_q^r.$$

The mechanism M must provide differential privacy to its input trajectories, which fundamentally means that the output of M should be insensitive to changes in its inputs. Differential privacy captures this notion by requiring that, whenever $\text{Adj}_B(u, \tilde{u})$ holds, the probability distributions of Mu and $M\tilde{u}$ satisfy

$$\mathbb{P}(Mu \in S) \leq e^\epsilon \mathbb{P}(M\tilde{u} \in S)$$

for all S in an appropriate σ -algebra.

Using a result from the literature, we now state a finite-time criterion which holds if and only if M keeps entire trajectories private. Below, the notation $v_{0:k} := (v(0), \dots, v(k))$ is used to refer to the first $k+1$ entries of $v \in \tilde{\ell}_q^r$.

Lemma 2: Let $\epsilon \geq 0$ be given. For a dynamical system, a mechanism M is ϵ -differentially private if and only if, for all u, \tilde{u} satisfying $\text{Adj}_B(u, \tilde{u}) = 1$ and for all times k ,

$$\mathbb{P}((M(u))_{0:k} \in A) \leq e^\epsilon \mathbb{P}((M(\tilde{u}))_{0:k} \in A) \text{ for all } A \in \mathcal{B}^{(k+1)r},$$

where \mathcal{B}^d is the Borel sigma-algebra on \mathbb{R}^d and r is the dimension of the output space.

Proof: See Lemma 2 in [15]. ■

In Lemma 2, the value of ϵ determines the level of privacy afforded to the input signals, and decreasing its value leads to improved privacy at the cost of adding higher variance noise. Typical values of ϵ in the literature range from 0.1 to $\ln 3$.

B. Joint Differential Privacy

We now elaborate on the application of joint differential privacy to Problem 1. To promote truth-telling behaviors, limits are imposed on the ability of any agent to reduce its cost by reporting a false state trajectory to the cloud. These limits are enforced using joint differential privacy, which is a relaxation of ordinary differential privacy for use in multi-agent systems, and it will be shown that this framework is sufficient for the goal of reducing an agent’s ability to benefit from misreporting its state. For joint differential privacy, the “system” of interest is comprised by the computations the cloud carries out in accordance with Update (1), and this point is discussed further below. For now it suffices to point out that the output of this system is a tuple of private forms of all q_i ’s. Denoting the

private form of $q_i(k)$ by $\tilde{q}_i(k)$ (whose exact form will be given later), the output of the cloud at time k is denoted by

$$y(k) = (\tilde{q}_1(k), \dots, \tilde{q}_N(k)).$$

Let \mathcal{M} denote a mechanism for joint differential privacy and let \mathcal{M}_{-i} denote the same mechanism with the i^{th} output removed, i.e., the output of \mathcal{M}_{-i} is

$$y_{-i}(k) = (\tilde{q}_1(k), \dots, \tilde{q}_{i-1}(k), \tilde{q}_{i+1}(k), \dots, \tilde{q}_N(k)).$$

A mechanism \mathcal{M} is joint differentially private if, for any $i \in [N]$, \mathcal{M}_{-i} preserves differential privacy for inputs adjacent with respect to i . Joint differential privacy for databases has been defined in [13], though, to our knowledge, it has not yet been used for dynamical systems. Using Lemma 2 in [15], the following lemma states a finite-time criterion for trajectory-level joint differential privacy.

Lemma 3: (Joint differential privacy for dynamical systems) Let the privacy parameter $\epsilon \geq 0$ be given and let \mathcal{M} be a mechanism whose output is an N -tuple. Then \mathcal{M} is ϵ -joint differentially private if and only if, for any $i \in [N]$, for all $u, \tilde{u} \in \tilde{\ell}_p^s$ satisfying $\text{Adj}_B^i(u, \tilde{u}) = 1$, all times k , and all $A \in \mathcal{B}^{(k+1)(n-n_i)}$, \mathcal{M} satisfies

$$\mathbb{P}((\mathcal{M}_{-i}(u))_{0:k} \in A) \leq e^\epsilon \mathbb{P}((\mathcal{M}_{-i}(\tilde{u}))_{0:k} \in A)$$

where \mathcal{B}^d is the Borel sigma-algebra on \mathbb{R}^d . ■

Joint differential privacy guarantees that when agent i 's input changes by a small amount, the outputs corresponding to other agents do not change by much. A useful characteristic of both ordinary and joint differential privacy is their resilience to post-processing, which guarantees that post-hoc transformations of private data cannot weaken the privacy guarantees afforded to that data. This result is formalized below.

Lemma 4: (Resilience to post-processing; [6, Proposition 2.1]) Let M be an ϵ -differentially private mechanism and let f be a function such that the composition $f \circ M$ is well-defined. Then $f \circ M$ is also ϵ -differentially private. ■

C. The Laplace Mechanism

To enforce differential privacy for a particular choice of ϵ , noise must be added somewhere in the system, and the distribution of this noise must be determined. One common mechanism in the literature draws noise from the Laplace distribution, used in both [7] and [15], and this mechanism is used to provide ϵ -joint differential privacy. To define this mechanism, the notion of the ℓ_p sensitivity of a system is now introduced.

Definition 2: The ℓ_p sensitivity of a deterministic, causal system \mathcal{G} is defined as

$$\Delta_p \mathcal{G} = \sup_{u, \tilde{u}: \text{Adj}_B(u, \tilde{u})=1} \|\mathcal{G}(u) - \mathcal{G}(\tilde{u})\|_{\ell_p}. \quad \blacklozenge$$

The Laplace mechanism is stated in terms of the ℓ_1 sensitivity of a system. Below, the notation $\text{Lap}(b)$ denotes a scalar Laplace distribution with mean zero and scale parameter b , i.e.,

$$\text{Lap}(b) := \frac{1}{2b} \exp\left(-\frac{|x|}{b}\right).$$

Lemma 5: (Laplace mechanism; [15, Theorem 4]) Let $\epsilon \geq 0$ be given. The Laplace mechanism defined by

$$M(u) = \mathcal{G}(u) + w$$

with $w(k) \sim \text{Lap}(b)^r$ is ϵ -differentially private for $b \geq \Delta_1 \mathcal{G}/\epsilon$, with r the dimension of the output space of the system. ■

Lemma 5 says that one can make a system private by adding noise drawn from a Laplace distribution to that system's output at each timestep. This idea is implemented for Problem 1 in the next section.

IV. OPTIMIZING UNDER JOINT DIFFERENTIAL PRIVACY

In this section, the privacy results in Section III are applied to Problem 1, and Algorithm 1 is made joint differentially private. Formally, g and g_{x_i} are treated as memoryless dynamical systems, and joint differential privacy is used to ensure that they keep the agents' state trajectories, which are the inputs to these systems, private.

A. Stochastic Optimization Algorithm

From Lemma 3, we see that enforcing joint differential privacy for the states in the network requires that the cloud make $q_i(k) := g_{x_i}(x(k))^T \mu(k)$ private before it is sent to agent i . To make $g_{x_i}(x(k))$ private, noise can be added to it directly and this is done below. To make $\mu(k)$ private, we use the fact that computing $\mu(k)$ relies on $g(x(k-1))$, and add noise to make $g(x(k-1))$ private. Then computing $\mu(k)$ is joint differentially private by the post-processing property in Lemma 4. Similarly, if the noisy forms of both $g_{x_i}(x(k))$ and $\mu(k)$ are joint differentially private, their product is as well, again by Lemma 4. Adding noise in this way, Algorithm 1 is modified to state the joint differentially private optimization algorithm below.

Algorithm 2:

Step 0: For all $i \in [N]$, initialize agent i with $x_i(0)$, X_i , f_i , $\{\alpha_k\}_{k \in \mathbb{N}}$, and $\{\gamma_k\}_{k \in \mathbb{N}}$. Initialize the cloud with \bar{x} , g , f_{lower} , $\{\alpha_k\}_{k \in \mathbb{N}}$, and $\{\gamma_k\}_{k \in \mathbb{N}}$. Let the cloud compute M before the system begins optimizing. Set $k = 0$.

Step 1: For all $i \in [N]$, the cloud computes

$\tilde{q}_i(k) := (g_{x_i}(x(k)) + w_i(k))^T \mu(k)$ and sends it to agent i .

Step 2: Agent i computes

$$x_i(k+1) = \Pi_{X_i} \left[x_i(k) - \gamma_k (\nabla_i f_i(x_i(k)) + \tilde{q}_i(k) + \alpha_k x_i(k)) \right],$$

and sends a state value $x_i(k+1)$ to the cloud.

Step 3: The cloud computes

$$\mu(k+1) = \Pi_M [\mu(k) + \gamma_k (L_\mu(k) + w_g(k) - \alpha_k \mu(k))].$$

Step 4: Set $k := k+1$ and return to Step 1. ◇

To solve Problem 1, Algorithm 2 must implement joint differential privacy using the noise terms w_g and w_i , while still converging to a minimum. The following theorem gives conditions on w_g and each w_i under which convergence to a minimum is guaranteed. Then Theorem 2 shows that joint differential privacy is achieved under these conditions, and Section V shows that each agent's incentive for misreport is indeed limited due to joint differential privacy.

Theorem 1: Let $(\hat{x}, \hat{\mu})$ denote the least-norm saddle point of L . Algorithm 2 satisfies

$$\lim_{k \rightarrow \infty} \mathbb{E}[\|x(k) - \hat{x}\|_2^2] = 0 \text{ and } \lim_{k \rightarrow \infty} \mathbb{E}[\|\mu(k) - \hat{\mu}\|_2^2] = 0$$

if i. $w_i(k)$ and $w_g(k)$ have zero mean for all k

ii. $\alpha_k = \bar{\alpha}k^{-c_1}$ and $\gamma_k = \bar{\gamma}k^{-c_2}$, where $0 < c_1 < c_2$, $c_1 + c_2 < 1$, $0 < \bar{\alpha}$, and $0 < \bar{\gamma}$.

Proof: See [18, Theorem 6]. ■

It remains to be shown that Condition i of Theorem 1 can be satisfied when joint differential privacy is implemented, and this is done next.

B. Calibrating Noise for Joint Differential Privacy

Here the systems being made private are g and g_{x_i} , and the mechanisms acting for joint differential privacy add noise to $g(x(k-1))$ when computing $\mu(k)$ and add noise to $g_{x_i}(x(k))$ when computing $\tilde{q}_i(k)$. It was shown in Section III that the noise added in Algorithm 2 will enforce ϵ -joint differential privacy as long as it has large enough variance. To determine the variance of noise that must be added by the Laplace mechanism, bounds are derived on the ℓ_1 sensitivity of each g_{x_i} and g below.

Lemma 6: For the relation Adj_B , the ℓ_1 sensitivities of g_{x_i} and g satisfy $\Delta_1 g_{x_i} \leq L_{g,i}B$ and $\Delta_1 g \leq K_g B$.

Proof: See [9]. ■

Using Lemmas 5 and 6, we see that if $w_i(k) \sim \text{Lap}(b_i)$ with $b_i \geq \Delta_1 g_{x_i}/\epsilon$ and $w_g(k) \sim \text{Lap}(b_g)$ with $b_g \geq \Delta_1 g/\epsilon$ for all k , then all states are afforded (ordinary) ϵ -differential privacy in Algorithm 2. It turns out that this privacy and its resilience to post-processing imply that ϵ -joint differential privacy holds as well, which is stated formally in the following theorem.

Theorem 2: Consider the mechanism \mathcal{M} defined by

$$\mathcal{M}(x(k)) = (\tilde{q}_1(k), \dots, \tilde{q}_N(k)),$$

with $\tilde{q}_i(k) := (g_{x_i}(x(k)) + w_i(k))^T \mu(k)$ as defined in Algorithm 2 and $\mu(k)$ computed as in Algorithm 2. At each time k , if $w_i(k) \sim \text{Lap}(b_i)$ with $b_i \geq \Delta_1 g_{x_i}/\epsilon$ and $w_g(k) \sim \text{Lap}(b_g)$ with $b_g \geq \Delta_1 g/\epsilon$, then \mathcal{M} is ϵ -joint differentially private.

Proof: We examine \mathcal{M}_{-i} for an arbitrary $i \in [N]$ whose output at time k is

$$\mathcal{M}_{-i}(x(k)) = (\tilde{q}_1(k), \dots, \tilde{q}_{i-1}(k), \tilde{q}_{i+1}(k), \dots, \tilde{q}_N(k)).$$

Examining some $j \in [N] \setminus \{i\}$, the j^{th} output $\mathcal{M}_j(x(k))$ is

$$\mathcal{M}_j(x(k)) = \tilde{q}_j(k) := (g_{x_j}(x(k)) + w_j(k))^T \mu(k).$$

In light of the fact that $w_g(k-1) \sim \text{Lap}(b_g)$ with $b_g \geq \Delta_1 g / \epsilon$, we see that $\mu(k)$ keeps $x(k-1)$ ϵ -differentially private. Similarly, $g_{x_j}(x(k)) + w_j(k)$ keeps $x(k)$ ϵ -differentially private because $w_j(k) \sim \text{Lap}(b_j)$ and $b_j \geq \Delta_1 g_{x_j} / \epsilon$. By Lemma 4, the product $\tilde{q}_j(k) = (g_{x_j}(x(k)) + w_j(k))^T \mu(k)$ keeps x ϵ -differentially private because it is the result of post-processing two differentially private quantities. By the same reasoning,

$$\mathcal{M}_{-i}(x(k)) = (\tilde{q}_1(k), \dots, \tilde{q}_{i-1}(k), \tilde{q}_{i+1}(k), \dots, \tilde{q}_N(k))$$

simply post-processes differentially private information. From Lemma 3 we conclude that \mathcal{M} is ϵ -joint differentially private. \blacksquare

The next section describes the application of this mechanism to inducing approximately-truthful behavior in multi-agent optimization through the computation of β -approximate minima.

V. COMPUTING β -APPROXIMATE MINIMA

This section presents the main result of the paper: joint differential privacy results in there being only minimal incentive for an agent to misreport its state to the cloud. Toward showing this result, a uniform upper bound on f_i over X_i is first presented.

Lemma 7: Let \bar{x} denote a Slater point for g . Then for all $i \in [N]$, $f_i(x_i) \leq \lambda_i$ for all $x_i \in X_i$, where $\lambda_i := f_i(\bar{x}_i) + K_i D_i$.

Proof: Using the Mean-Value Theorem and the Cauchy-Schwarz inequality, we have

$$\begin{aligned} f_i(x_i) &= f_i(\bar{x}_i) + \nabla f_i(z_i)^T (x_i - \bar{x}_i) \\ &\leq f_i(\bar{x}_i) + \|\nabla f_i(z_i)\| \cdot \|x_i - \bar{x}_i\|, \end{aligned}$$

for some $z_i \in X_i$. The result follows by bounding $\|\nabla f_i(z_i)\|$ by K_i and bounding $\|x_i - \bar{x}_i\|$ by D_i . \blacksquare

The next lemma bounds the difference in f_i at points along two feasible state trajectories for agent i .

Lemma 8: For any $x_i^1, x_i^2 \in \tilde{\ell}_{p_i}^{s_i}$ and any time k , one finds

$$|f_i(x_i^1(k)) - f_i(x_i^2(k))| \leq \rho_i := \min\{K_i D_i, 2\lambda_i\}.$$

Proof: Using the Lipschitz property of f_i ,

$$|f(x_i^1(k)) - f(x_i^2(k))| \leq K_i \|x_i^1(k) - x_i^2(k)\|_1 \leq K_i D_i.$$

On the other hand, the triangle inequality gives

$$|f(x_i^1(k)) - f(x_i^2(k))| \leq |f(x_i^1(k))| + |f(x_i^2(k))| \leq 2\lambda_i,$$

where the second inequality follows from Lemma 7. ■

The main result of the paper is now presented. Below, each expected value is over the randomness introduced by the mechanism \mathcal{M} . For clarity, this theorem tracks the state agent i has reported to the cloud and the state trajectory of every *other* agent. The notation $\mathbb{E}[f_i(x_i(k))|y_i, v_{-i}]$ is used to denote agent i 's cost at time k when agent i has reported the trajectory y_i to the cloud and every other agent has followed the trajectory v . The symbol x_i is always used as the argument to f_i because f_i always depends on the true state of agent i , not the state it reports.

Theorem 3: Suppose Assumptions 1 and 2 hold. Let the agents and cloud execute Algorithm 2 with the cloud implementing the mechanism \mathcal{M} . Then for $\epsilon \in (0, 1)$, all agents sharing their true states in Algorithm 2 results in a β -approximate minimum. In particular, at all times k and for any state trajectories $x, \tilde{x} \in \tilde{\ell}_p^s$ satisfying $\text{Adj}_B^i(x, \tilde{x}) = 1$, we have

$$\mathbb{E}[f_i(x_i(k))|x_i, x_{-i}] \leq \mathbb{E}[f_i(x_i(k))|x'_i, \tilde{x}_{-i}] + \beta,$$

where $\beta = 2 \max_{i \in [N]} \rho_i + \epsilon \lambda_i$, and where x'_i is a misreported state trajectory from agent i .

Proof: From Lemma 8 we have

$$\mathbb{E}[f_i(x_i(k))|x_i, x_{-i}] \leq \mathbb{E}[f_i(x_i(k))|\tilde{x}_i, x_{-i}] + \rho_i. \quad (19)$$

Using Theorem 2 and the definition of ϵ -joint differential privacy we find

$$\mathbb{E}[f_i(x_i(k))|\tilde{x}_i, x_{-i}] \leq e^\epsilon \mathbb{E}[f_i(x_i(k))|\tilde{x}_i, \tilde{x}_{-i}],$$

which we substitute into Equation (V) to find

$$\mathbb{E}[f_i(x_i(k))|x_i, x_{-i}] \leq e^\epsilon \mathbb{E}[f_i(x_i(k))|\tilde{x}_i, \tilde{x}_{-i}] + \rho_i.$$

Using $e^\epsilon \leq 1 + 2\epsilon$ for $\epsilon \in (0, 1)$ gives

$$\mathbb{E}[f_i(x_i(k))|x_i, x_{-i}] \leq \mathbb{E}[f_i(x_i(k))|\tilde{x}_i, \tilde{x}_{-i}] + 2\epsilon \mathbb{E}[f_i(x_i(k))|\tilde{x}_i, \tilde{x}_{-i}] + \rho_i, \quad (20)$$

where we apply Lemma 7 to get

$$\mathbb{E}[f_i(x_i(k))|x_i, x_{-i}] \leq \mathbb{E}[f_i(x_i(k))|\tilde{x}_i, \tilde{x}_{-i}] + 2\epsilon \lambda_i + \rho_i. \quad (21)$$

A second application of Lemma 8 gives

$$\mathbb{E}[f_i(x_i(k))|\tilde{x}_i, \tilde{x}_{-i}] \leq \mathbb{E}[f_i(x_i(k))|x'_i, \tilde{x}_{-i}] + \rho_i,$$

and substituting this inequality into Equation (V) gives

$$\mathbb{E}[f_i(x_i(k))|x_i, x_{-i}] \leq \mathbb{E}[f_i(x_i(k))|x'_i, \tilde{x}_{-i}] + \beta,$$

as desired. ■

While the ρ_i term in β is a feature of the problem itself, the $\epsilon \lambda_i$ term results directly from the untruthfulness of agent i , and it is precisely this term which can be influenced using the privacy parameter ϵ , allowing a network operator to directly counteract the influence of false information. Of course, shrinking ϵ requires that more noise




Fig. 2. A plot of the distance to the saddle point in Problem 0 in the primal space (lower curve) and dual space (upper curve) when all agents are truthful.

be added which, in general, degrades performance in the system. One must therefore balance the two objectives of incentivizing truthful information sharing and system performance based upon the needs in a particular application.

VI. SIMULATION RESULTS

A simulation was run consisting of $N = 8$ agents each with $x_i \in \mathbb{R}^2$, and $m = 4$ constraints. The constraint $X_i = [-10, 10]^2$ for all $i \in [N]$, and $f_i(x_i) = \frac{1}{2}\|x_i - t_i\|_2^2$, for all $i \in [N]$; the value of each t_i can be found in Table I. The constraints used were

$$g(x) = \begin{pmatrix} \|x_1 - x_2\|_2^2 + \|x_1 - x_3\|_2^2 - 5 \\ \|x_4 - x_5\|_2^2 + \|x_4 - x_6\|_2^2 - 3 \\ \|x_7 - x_8\|_2^2 + \|x_7 - x_6\|_2^2 - 3 \\ \|x_5 - x_3\|_2^2 + \|x_5 - x_7\|_2^2 - 5 \end{pmatrix}.$$

The value $B = 3$ was used for adjacency and the privacy parameter was chosen to be $\epsilon = \ln 3$. The distributions of noise added are shown in Table I; in addition to those values, $w_g \sim \text{Lap}(327.69)$. The stepsize and regularization parameters were chosen to be $\gamma_k = 0.01k^{-3/5}$ and $\alpha_k = 0.5k^{-1/3}$. Two adjacent problems were run for 250,000 timesteps each, with agent 6 being untruthful in one of them. In that run, agent 6 reported its unconstrained minimizer t_6 to the cloud at each timestep instead of its actual state.

Figure 2 plots the distance to the saddle point in the primal and dual spaces when all agents are truthful. The final error values were $\|x(250,000) - \hat{x}\|_2 = 0.5367$ and $\|\mu(250,000) - \hat{\mu}\|_2 = 0.6870$, indicating close convergence in each space despite the large amount of noise in the system.

Figure 3 plots the decrease in cost agent 6 sees in misreporting its state. This decrease is not only bounded by β , but is bounded above by 0.1β for all time, indicating that β may be loose in some cases; this is not surprising given that β uses Lipschitz constants and set diameters, which are “worst-case” in the sense that they give maximum values over all possible states. Nonetheless, the algorithm can be seen both to converge and compute a β -approximate minimum, indicating that Problem 1 has been solved.




Fig. 3. A plot of the decrease in cost agent 6 attains through misreporting its state. The ordinate is normalized by β .

i	t_i	Distribution of w_i
1	$[6 \ -4]^T$	Lap(10.92)
2	$[2 \ 2]^T$	Lap(5.46)
3	$[-7 \ 7]^T$	Lap(5.46)
4	$[8 \ -9]^T$	Lap(10.92)
5	$[3 \ -7]^T$	Lap(16.38)
6	$[10 \ 10]^T$	Lap(10.92)
7	$[-10 \ -10]^T$	Lap(16.38)
8	$[6 \ -6]^T$	Lap(5.46)

TABLE I

VALUES OF t_i AND DISTRIBUTIONS OF w_i FOR $i \in [8]$.

VII. CONCLUSION

It was shown that joint differential privacy can be used in multi-agent optimization to incentivize truthful information sharing. Applications of the work presented here include any multi-agent setting in which the iterates of an optimization algorithm correspond to some physical quantity of interest. Future directions include allowing asynchronous communications in order to account for systems with communication latency and poor channel quality.

REFERENCES

- [1] AB Bakushinskii and BT Polyak. Solution of variational inequalities. *Doklady Akademii SSSR*, 219(5):1038–1041, 1974.
- [2] Dimitri P Bertsekas, Angelia Nedic, and Asuman E Ozdaglar. *Convex analysis and optimization*. Athena Scientific, 2003.
- [3] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends Mach. Learn.*, 3(1), January 2011.
- [4] S. Braynov and M. Jadliwala. Detecting malicious groups of agents. In *Multi-Agent Security and Survivability, 2004 IEEE First Symposium on*, pages 90–99, Aug 2004.
- [5] Jorge Cortes, Sonia Martinez, Timur Karatas, and Francesco Bullo. Coverage control for mobile sensing networks. In *Robotics and Automation, 2002. Proceedings. ICRA'02. IEEE International Conference on*, volume 2, pages 1327–1332. IEEE, 2002.
- [6] C. Dwork and A. Roth. The algorithmic foundations of differential privacy. *Theoretical Computer Science*, 9(3-4):211–407, 2013.

- [7] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Proc. of the Third Conference on Theory of Cryptography*, TCC'06, pages 265–284, Berlin, Heidelberg, 2006. Springer-Verlag.
- [8] A. Fagiolini, M. Pellinacci, G. Valenti, G. Dini, and A. Bicchi. Consensus-based distributed intrusion detection for multi-robot systems. In *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on*, pages 120–127, May 2008.
- [9] M.T. Hale and M. Egerstedt. Cloud-enabled multi-agent optimization with constraints and differentially private states. 2016. Submitted for publication. Available at <http://arxiv.org/abs/1507.04371>.
- [10] S. Han, U. Topcu, and G. J. Pappas. Differentially private distributed constrained optimization. *IEEE Transactions on Automatic Control*, PP(99):1–1, 2016.
- [11] Shuo Han, U. Topcu, and G.J. Pappas. An approximately truthful mechanism for electric vehicle charging via joint differential privacy. In *American Control Conference (ACC), 2015*, July 2015.
- [12] Justin Hsu, Zhiyi Huang, Aaron Roth, and Zhiwei Steven Wu. Jointly private convex programming. In *Proc. of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '16, 2016.
- [13] M. Kearns, M. Pai, A. Roth, and J. Ullman. Mechanism design in large games: Incentives and privacy. *CoRR*, abs/1207.4084, 2012.
- [14] H. W. Kuhn and A. W. Tucker. Nonlinear programming. In *Proc. of the 2nd Berkeley Symposium on Math. Stat. and Prob.*, Berkeley, Calif., 1951. University of California Press.
- [15] J. Le Ny and G.J. Pappas. Differentially private filtering. *Automatic Control, IEEE Transactions on*, 59(2):341–354, Feb 2014.
- [16] Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '07, pages 94–103, Washington, DC, USA, 2007. IEEE Computer Society.
- [17] M.H. Nazari, Z. Costello, M.J. Feizollahi, S. Grijalva, and M. Egerstedt. Distributed frequency control of prosumer-based electric energy systems. *Power Systems, IEEE Transactions on*, 29(6), Nov 2014.
- [18] BT Poljak. Nonlinear programming methods in the presence of noise. *Math. programming*, 14(1):87–97, 1978.
- [19] Daniel E Soltero, Mac Schwager, and Daniela Rus. Decentralized path planning for coverage tasks using gradient descent adaptive control. *The International Journal of Robotics Research*, 2013.